ESD-TDR-64-449

DECISION MAKING WITH FREE OPERANT RESPONSES

TECHNICAL DOCUMENTARY REPORT NO. ESD-TDR-64-449

JUNE 1964

William C. Holz

DECISION SCIENCES LABORATORY
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
L. G. Hanscom Field, Bedford, Massachusetts

Project 7682, Task 768204

(Prepared under Contract No. AF 19 (628)-2404 by the Harvard University Committee on Programmed Instruction, Cambridge, Massachusetts.)

DDC AVAILABILITY NOTICES

ESD-TDR-64-449

DECISION MAKING WITH FREE OPERANT RESPONSES

TECHNICAL DOCUMENTARY REPORT NO. ESD-TDR-64-449

JUNE 1964

William C. Holz

DECISION SCIENCES LABORATORY
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
L.G. Hanscom Field, Bedford, Massachusetts

Project 7682, Task 768204

(Prepared under Contract No. AF 19 (628)-2404 by the Harvard University Committee on Programmed Instruction, Cambridge, Massachusetts.)

FOREWORD

This report covers a portion of the applied research program of the Decision Sciences Laboratory. The research was conducted under Contract AF 19(628)-2404 in support of Project 7682, Man-Computer Information Processing, Task 768204, Automated Training for Information Systems.

Dr. William C. Holz of Harvard University was the principal investigator and author. Dr. Sylvia R. Mayer of the Decision Sciences Laboratory was the Air Force technical monitor.

The research was conducted in facilities provided by the Harvard Committee on Programmed Instruction which is supported by the Carnegie Corporation.

ABSTRACT

These experiments explored the suitability of free operant techniques in the investigation of choice behavior and decision making. Young adults were the subjects. Two response manipulanda were available; and points were intermittently scheduled in different proportions for each. The number of points at the end of the session determined the subjects' payment. The schedule by which the points could result was the independent variable; and the relative frequency of the two responses, which represented the subject's choice, was the dependent variable. When the points were scheduled randomly in time, the anticipated result on the basis of previous findings was that the relative frequency of response would match the relative frequency of points. The observed result did not clearly follow this pattern. Over the period studied, the pattern was one of approximately equal responding to both choices regardless of the relative frequency of points obtained.

In two similar experiments the points were scheduled randomly in time, but a requirement was added that responses must be spaced at two second intervals to produce a point. The purpose of this experiment was to determine if reducing the high rate of response observed in the previous experiment would lead the relative frequency of response to conform with the expected pattern. Under these conditions, the results closely approximated the matching model. Further, as the relative frequency of reinforcement for the two responses was changed, the relative frequency of these responses followed directly. With these modifications of the schedule, the results show continuity with previous findings, and indicate that the probability of action is a direct function of the probability of reinforcement.

REVIEW AND APPROVAL

This technical documentary report has been reviewed and is approved.

FOR THE COMMANDER

Joseph T. Begley

JOSEPH T. BEGLEY
Chief, Applications Division
Decision Sciences Laboratory

Roy Morgan

ROY MORGAN
Colonel, USAF
Director, Decision Sciences Laboratory

KEY WORD LIST

1. DECISION MAKING
2. BEHAVIOR
3. REACTION (PSYCHOLOGY)
4. HUMAN ENGINEERING
5. LEARNING
6. CONDITIONED REFLEX
7. EXPERIMENTAL DATA
8. PROGRAMMED INSTRUCTION

TABLE OF CONTENTS

SECTION 1

INTRODUCTION

The topics of "decision making" and "choice behavior" have held considerable interest for psychologists in recent years. Broadly stated, these topics entail predicting the behavior of a person confronted by an uncertain situation. Two or more courses of action are open, but the outcome of each is uncertain. The problem is to identify the significant variables which influence his choice or decision.

Much of the research on these topics has been aimed directly at the problem of human performance in situations which closely parallel those of everyday life. Games of chance, for example, have frequently served as the experimental situation. The subject plays the game against the experimenter who sets up different odds, risks and the like, and records the subject's response to them. These investigations have been primarily concerned with fitting mathematical models to the performance, and these models provide us with theories of broad generality for predicting such behavior on a probabilistic basis. Although the situations to which the models apply may take a number of forms, the essential elements would seem to be encompassed by the description in the first paragraph.

When the problem is stated in its general form, there is yet another area of research which deals with very similar issues. This is the traditional area of learning and conditioning. The rat in the T-maze is faced with the choice of right and left turns, and the monkey must decide when to press the lever. However, the findings from this area are typically not integrated with those of the former. Perhaps initially they were not incorporated because conditioning and learning experiments were limited to investigations of consistent rather than uncertain situations. The behavior of the rat, for example, was studied when turns to a particular side always led to the food pellet. And, later, when intermittent schedules of reinforcement were introduced, only one response was studied. Recent developments in learning experiments, however, have broadened their scope to include situations which very closely parallel the "decision making" paradigm. As a result, findings from this area take on new importance. While the first area attempts to provide broad predictive generalizations, the latter area offers the potential of experimentally isolating the underlying processes. The resulting knowledge could thereby assist in attempts to control the process, as well as to predict its course.

In this regard, the work of Herrnstein is notable. He studied the behavior of pigeons when two responses were concurrently reinforced on intermittent schedules (Herrnstein, 1961). These studies utilized the free operant response in the context of complex stimulus and multiple response conditions. A basic discovery he has made is that rate of response is directly related to the rate of reinforcement. Previous studies with a single response procedure had at best shown a monotonic relationship between reinforcement frequency and response frequency which was clearly non-linear. A direct relationship between reinforcement and response had been postulated by theorists (Skinner, 1938) but until these investigations it had never been demonstrated experimentally.

Besides advancing our understanding of how reinforcement operates, his

findings have considerable potential for a greater understanding for the specific problems of decision making. Consider, for example, another experiment (Herrnstein, 1964a). A complex schedule was used in which alternative courses of action led to either a fixed interval or a variable interval schedule of primary reinforcement. The pigeon showed marked preference for the course leading to the variable interval schedule. This was true in spite of the fact that a lower overall rate of reinforcement resulted because of this. This seems analogous to the gambler, who continues to play in spite of obviously negative odds. From the point of view of a rational theory, such results are enigmatic; yet as Herrnstein points out, they seemingly occur as a lawful function of the effects of reinforcement. As such, further analysis of such effects of reinforcement seems important for advances in theory.

In another experiment (Herrnstein, 1964b) the alternative courses of action led to either a variable ratio or a variable interval schedule of positive reinforcement. By careful manipulation of the schedule values (i.e., the mean number of responses required for reinforcement and the mean interreinforcement interval), he found that rate of reinforcement per unit time, rather than rate of reinforcement per number of responses, was the critical factor in determining choice. In experiments using trials, the two variables of rate according to time and according to number are inextricably entwined. Only by using the free operant procedure in the context of the concurrent chain schedule, could the effects of these two variables be separated. These results suggest that the procedures developed in the study of the free operant response will be useful analytic tools.

Thus, with the extension of learning experiments to more complex types of schedules, the results and procedures of this field of experimentation would seem to have considerable potential for investigations of decision making. The emphasis here is on an analysis of the effects of reinforcement, utilizing a procedure which permits assessment of the absolute as well as the relative probabilities of response. The purpose of the experiments to be reported is to assess the suitability of applying these techniques directly with human subjects, and thereby to determine the generality of some of the findings with lower organisms. Of particular relevance in this regard is the work of Herrnstein (1961) which demonstrated that when two responses are concurrently reinforced, the relative frequency of the two responses is in proportion to the relative frequency of reinforcement received.

The experiments to be reported employed a button pressing response with adult human subjects. The button pressing response was selected because it is arbitrary with respect to the consequences, it is unambiguously defined by an electrical circuit, it can occur over a wide range of rates, and it would not seem to interact in any significant way with the particular history of the experiences of the subjects. The simple arbitrary response seemed most suitable for the analysis before extending the generalizations to the more complex repertoire of the adult human. Similarly, points exchanged for money seemed to be the most suitable reinforcer. This allowed precise determination of the relationship of the reinforcing stimulus to the behavior of the subject. The points could be delivered immediately at the prescribed times and did not require the presence of another person in the experimental environment.

In general, then, the subject was faced with a situation in which one alternative was more favorable than the other. The major independent variable was the frequency with which a response was reinforced. Variable interval schedules of

positive reinforcement associated with each response allowed the experimenter to manipulate the potential relative frequencies for the two responses and yet maintain a random pattern. The fact that the randomness of the schedule was dependent on time, and not upon the subject's behavior (as it would be with a variable ratio schedule), meant that the subject's absolute rate of response was free to vary over a wide range and still yield the same relative frequencies of reinforcement. Thus, responding was not forced to match the frequency of reinforcement as a requirement of the schedule. Variable interval schedules with small mean interreinforcement times were selected so that the differences between the schedules could exert their effect quickly. The situation was thus designed to be comparable in major outline with the experiments conducted by Herrnstein and to maintain the basic paradigm of the studies on decision making with human subjects.

SECTION 2

METHOD

Subjects: Thirty-seven male and female college students were studied throughout the course of these experiments. The subjects were secured through an advertisement in the Harvard newspaper and a notice in the student employment service office. The subjects were in their late teens or early twenties. They received either a flat fee of $1.50 per hour, or 1¢ for every point earned, whichever was greater.

Apparatus: During the experimental sessions, the subject was isolated in a 10 x 15 foot room. An 18 inch window fan provided ventilation for the room. A strip of metal vibrated against the revolving shaft of the fan to provide a masking noise for sounds from the control apparatus.

The subject was seated before an intelligence panel which contained three push buttons, several display lights, and a digital counter. Two of the push buttons were for recording the subjects' responses to the experimental conditions. These buttons, located on the front of the panel, were 3/4 inch in diameter and were separated 11 inches center to center. A third button was located on the upper left side of the panel. The experimental procedures required that the subject keep this third button depressed during the experimental periods in order that the experiment continue. The arrangement of the buttons and the necessity that one hand be occupied with the third button prevented the subject from responding on the two buttons simultaneously. A central red light indicated to the subject that the experimental session was in progress and was extinguished at the end of each experimental period. Two green lights, located in the upper portion of the panel on either side of the digital counter, flashed every time the counter advanced. These constituted the reinforcing stimulus.

The intelligence panel was connected to control equipment located in a separate room. This equipment consisted of standard electromagnetic devices. A response was recorded each time the subject depressed one of the response buttons during an experimental session. Impulse counters and a cumulative response recorder collected these data. After varying intervals of time, a response activated the counter on the intelligence panel. The schedule contingencies which determined when responses would activate the counter are described in the procedure section. For all contingencies responses on a button did not activate the counter until .5 sec. or more had elapsed after the subject initiated a change to that button. This changeover delay (COD) was introduced to minimize superstitious contingencies of reinforcement associated with changing keys (see, for example, Herrnstein, 1961; Catania & Cutts, 1963).

The electrical circuits which determined when responses on either button would produce points were independent of one another. The basic unit of these circuits was a "tape" timer. Holes punched in 8 mm. film tape determined the minimum time intervals between successive points. The distribution of the time intervals was irregular with the restrictions that intervals were equalized over 5 min. periods and had a minimum duration of 3 sec. Several different mean time intervals were used during the experiments and they are described in the procedure section.

Procedure: A laboratory assistant gave the subjects standard instructions at the beginning of the experiment. These instructions were as follows:

> You are to press the buttons on the panel. Occasionally, your press will advance the counter. Your task is to maximize this count. The amount of money earned is proportionate to the number of counts accumulated.
>
> Your time will be divided into ten 5 minute sessions with a one minute interim between each session. The experiment begins with a one minute break before the first session. The working period will begin when the red light in the center of the panel is illuminated. You must hold in this button (points) for the equipment to work. Try to earn the maximum amount of money during each five minute period. This amount could be as high as thirty cents per five minutes.
>
> When the red light goes out, you may stop. Record the number appearing on the counter on the sheet of paper, and reset the counter by depressing the small lever. The next period will begin shortly.
>
> You may smoke if you like, but do not leave the room until someone comes for you.
>
> Summary:
> 1. Start when the center red light is illuminated.
> 2. Stop when the red light extinguishes.
> 3. Record numbers on counter -- reset counter.
> 4. Hold in button on the left side of panel while you are working.

The experiment started after the assistant left the room. The red light on the subject's control panel illuminated to indicate that the apparatus was ready. The subject then depressed the side button to start the equipment. Only during the period that the subject depressed the side button was the equipment active. After 5 min. the red light went off and remained off for 1 min. The red light then came on for another 5 min., and so forth. Thus, the daily session for a subject consisted of 10 such 5 min. samples.

The nomenclature provided by Ferster and Skinner (1957) will be used to describe the schedule contingencies. "Concurrent" (concur.) designates two schedules that are in force over the same period of time. In all of the experiments, the timing circuits which determined the availability of reinforcers for both keys were active throughout the experimental periods, and hence, the schedules were all concurrent. "Variable interval" (VI) refers to the fact that reinforcements were scheduled at irregular periods of time. By convention, the mean of these intervals is specified in minutes. In some of the experiments a "differential reinforcement of low rate" (drl) contingency was employed. With the drl, reinforcements are programmed according to the variable interval schedule and are delivered only to a response which is separated by a specified period of time from the preceding response. This period is specified in seconds.

Further, for all of the experiments, the contingency for the left button is given before that for the right button. For example: concur. VI .2 drl 2, VI 1 drl 2, designates that the left response will be reinforced at irregular periods of time averaging 12 sec. (.2 min.), provided that the reinforced response is separated from the preceding response by at least 2 sec.; and the right response will be reinforced at irregular periods averaging 1 min. when the responses are spaced at least 2 sec. apart.

Experiment I

For the first experiment points were delivered for responses according to two variable interval schedules of reinforcement, concur. VI .2 VI 1. Since the mean interval was .2 min. for the left response and 1 min. for the right response, the subject received an average of 5 reinforcements on the left key to 1 reinforcement on the right. The relative frequency of reinforcement on the right key was .17.

Experiment II

In the second experiment, a drl contingency was added stipulating that only responses spaced 2 sec. from the preceding response would be reinforced. The variable interval contingency remained in effect. The two schedules studied were concur. VI .2 drl 2 VI 1 drl 2, and concur. VI .5 drl 2 VI .25 drl 2.

Experiment III

In the third experiment, as in the second, the reinforcement schedule was concur. VI drl, VI drl; but, the mean values of the variable interval schedule varied. Thus, the schedules studied were: concur. VI .2 drl 2, VI 1 drl 2; concur. VI .25 drl 2, VI .5 drl 2; concur. VI .33 drl 2, VI .33 drl 2; concur. VI .5 drl 2, VI .25 drl 2; concur. VI 1 drl 2, VI .2 drl 2. Under optimal conditions, these schedules provide relative frequencies of reinforcement of .83:.17; .67:.33; .50:.50; .33:.67; .17:.83 on the left and right buttons.

SECTION 3

RESULTS

Experiment I

The scheduling of reinforcement according to the concur. VI .2, VI 1 led to high rates of response on both buttons. Response rates on the two buttons combined ranged from 102 to 274 responses per minute with a median of 221 for 14 subjects. Since high rates of responding were maintained on both buttons, the subjects collected nearly all of the reinforcement allocated by both VI schedules. The relative frequency of reinforcement, therefore, closely approximated the scheduled values of .83 for the left response and .17 for the right response.

Figure 1 shows the extent to which the observed relative frequency of response on one button (the right) departed from the relative frequency of reininforcement of that response. The abscissa of Fig. 1 shows the percent of the total responses emitted on the right button minus the percent of total reinforcement received there. Thus, if 20% of the total responses were on the right key and 17% of the total reinforcements were delivered to right key responses, the value would be 20 - 17 = 3. These values were then grouped into class intervals of 10. The ordinate for the figure indicates the number of subjects whose "scores" fell in the particular class interval. As was noted in the introduction, experiments have suggested that the relative frequency of response should equal the relative frequency of reinforcement. Hence, one would anticipate a zero difference. The observed differences ranged from 9 to 66 percent, with a median of 34. It will be noted, also, that since the abscissa values are all positive, the deviation from the expected pattern was in the direction of greater responding on the button with lower reinforcement frequency. Responding on the two keys tended toward a .50 - .50 split, regardless of the frequency of reinforcement. An equal distribution of the responses on both buttons (i.e., 50% on the right), while 17% of the reinforcements resulted on the right, would give a difference of 50 - 17 = 33. This closely approximates the observed results.

These data represent the performance of subjects who had been exposed to the contingencies for only one hour. The experiments with lower organisms found it necessary to provide much longer periods of exposure before the data reached a stable performance at the expected level. It was anticipated that the human subject would come under control of the relative frequencies of reinforcement more quickly than the animal subjects. This did not prove to be the case. Some subjects were studied for longer periods of time (2 to 6 hours); and, in general, there was a drift toward the expected values. However, considerable difficulty was experienced in obtaining subjects who would continue long enough to thoroughly test the procedure. Continued exposure seemed unfeasible from a practical standpoint.

Other investigators (e.g., Azrin, 1958; Weiner, 1962), who have studied the application of operant conditioning procedures to human learning, have reported similar deviations from expected performance. In their studies of fixed interval reinforcement, for example, they observed that the most common

departure from the fixed interval scallops was a linear rate of response throughout the interval. But, when the effort of the response was increased, or when all responses were penalized, the typical fixed interval scallops appeared. Thus, when the effort required for the response was increased and the overall rate of response was lowered, the reduction was selective. The result was a closer approximation to the general findings of conditioning experiments with other organisms.

The second experiment to be reported approached the problem of high rates in another way. A contingency was added which specified that only responses spaced a minimum interval from the preceding response could be reinforced. This differential reinforcement of low rates (drl) procedure was investigated in experiment II.

Experiment II

Figure 2 shows the departure of the relative frequency of response from the relative frequency of reinforcement, when the drl contingency was added to the same reinforcement schedule used for Experiment I. Similarly, Fig. 3 shows this departure for a second concurrent schedule with variable interval schedules of different mean intervals. It will be noted that the drl contingency resulted in much closer approximations to the expected zero difference. As in Fig. 1, these data represent the performance at the end of a single hour's exposure to the schedule.

When Figs. 2 and 3 are compared, the effect of the different mean values of the variable interval schedule are apparent. With the concur. VI .5 drl 2, VI .25 drl 2, the left button was the one associated with the lower frequency of reinforcement. With the concur. VI .2 drl 2, VI 1 drl 2 the low frequency button was the right. In both cases, the deviations are in the direction of overresponding on the button with the lower frequency of reinforcement. This result suggests that the biasing is not simply one of position. That is, a preference for the right key, which is not counteracted by the high frequency of reinforcement on the left, is not responsible for the overresponding. Rather, the biasing is toward the button with the low frequency of reinforcement. This tendency is minimized, but not eliminated, by the drl contingency.

Although the drl contingency brought greater consistency to the performance, it also produced difficulties. Of 22 subjects who were started with the drl procedure in effect, only 15 (68%) came under control of the schedule. That is, the other subjects did not space their responses sufficiently to obtain at least 25% of the possible reinforcements. It will be noted that none of the subjects were told the nature of the reinforcement contingencies. We did not want to risk the possibility of biasing the data by giving such instructions, but wanted to see if the schedule would, in and of itself, generate the anticipated performance. Such instructions, however, might be useful in speeding the acquisition of schedule control.

Another 20% of these subjects proved to be unsuitable because they extinguished on the button with the low frequency of reinforcement. As the drl contingency gradually came to control a low rate, reinforcements were more often obtained on the button with the variable interval schedule of higher reinforcement density. The absence of reinforcement on the other key led to extinction by the time drl control was fully established. For this reason, many of the subjects were started with the concur. VI .5 drl 2, VI .25 drl 2 which provided

FIG. 1.

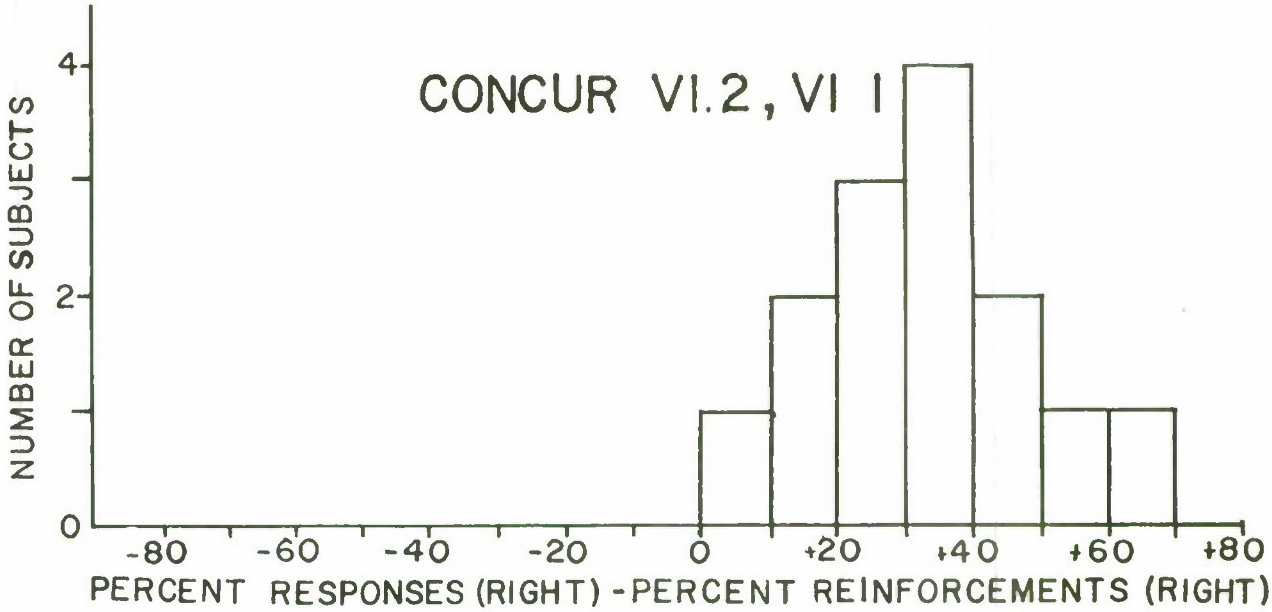

FIG. 2.

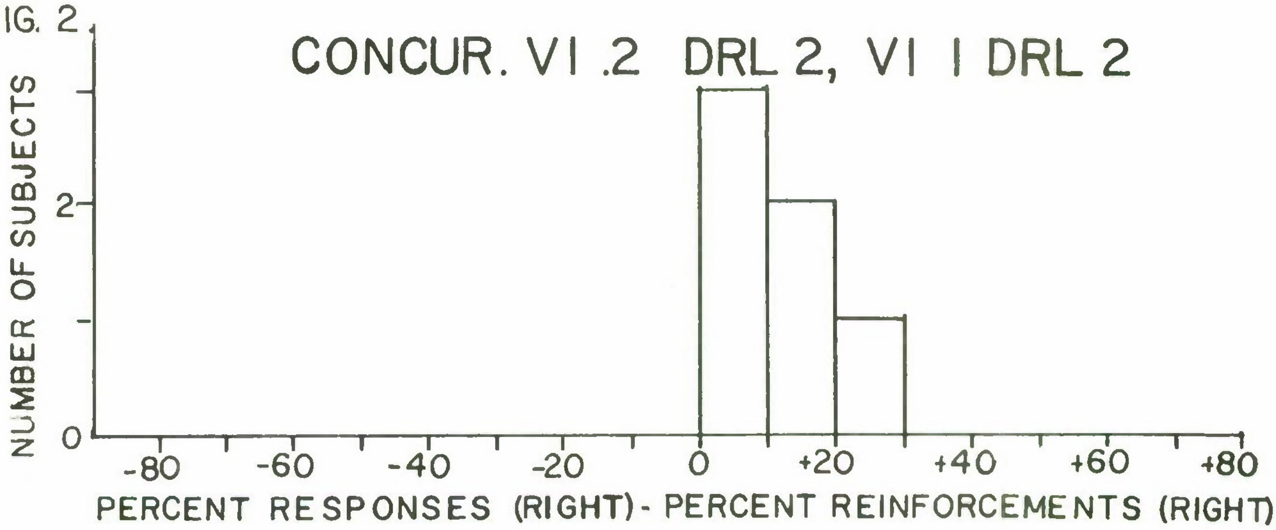

FIG. 3.

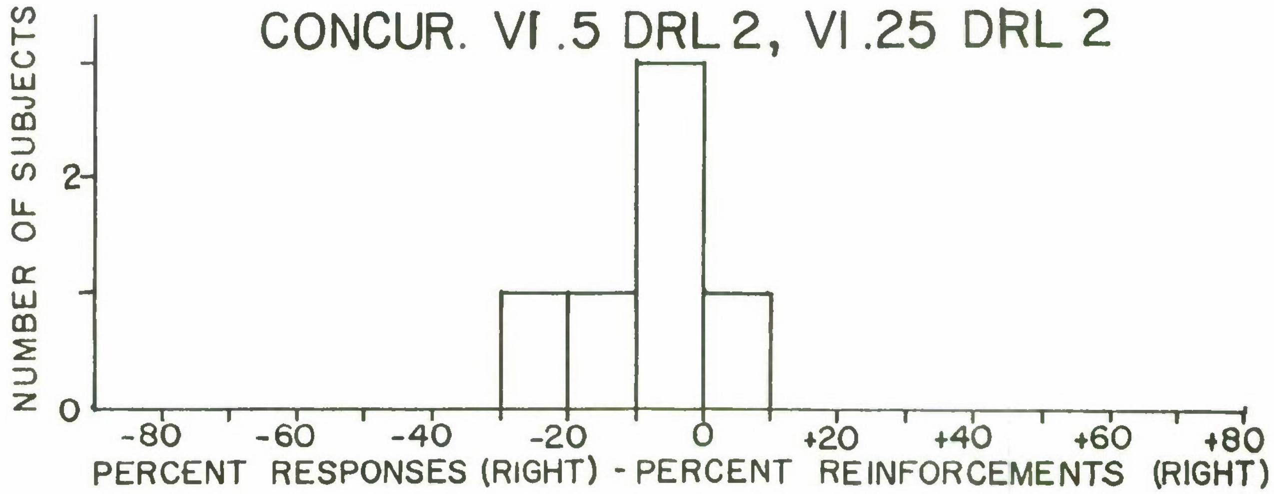

Fig. 1, 2, and 3. Differences between percent of total responses on the right button and the percent of the total reinforcements on that button. Note different reinforcement schedules used in each case.

a more equal density of reinforcement on the two keys. When extinction occurred with one response, 100% of the responses and 100% of the reinforcements occurred with the alternative response. This conforms to the expected relationship (i.e., the percent of responses equals the percent of reinforcements); but since this result is basically trivial these data have been omitted.

Similar problems prevented a more direct test of the efficiency of the drl schedule in minimizing deviations of relative frequencies of response and reinforcement. The drl contingency was added to the concur. VI .2, VI 1 schedule for 5 of the subjects. Only 2 subjects responded under control of the schedule after 2 to 4 hours exposure. The previous variable interval reinforcement seemed to greatly retard development of the drl control, and made comparison unfeasible (see also, Weiner, 1964). Conversely, when subjects were initially studied with the drl contingency in effect, removal of the contingency did not result in the typical variable interval schedule performance. Once the rate was lowered by the drl, the subjects continued to space their responses. The subjects, therefore, did not discriminate the removal of this contingency and continued to respond at the same low rate. An ABA, BAB, design, therefore, could not be accomplished.

Experiment III

In this experiment, the mean values of the variable interval schedules were varied. 14 subjects were studied on concur. VI drl, VI drl schedules with two or more different VI schedule values. In all cases, the total possible reinforcements for both responses was held constant while only the relative frequency was manipulated. The schedules used are specified in the procedure section. The subjects were roughly counterbalanced with respect to the order of the schedules. Since a number of the subjects discontinued before completion of the sequence, the counterbalancing was not precise; but, no order effect was apparent.

Fig. 4 shows the combined data of all subjects. In this figure, the ordinate represents the percent of responses on the right key; the abscissa, the percent of reinforcements on this same key. The points at which the percent of response equals the percent of reinforcement are represented by the solid line with the slope of 1. The data points represent the medians of the final three 5 min. periods, for each session after the subjects' performance had stabilized. In general, two stable sessions were observed before another schedule value was introduced. Replications with the same subjects are included in these data. It will be noted that the points do not line up immediately above the abscissa points associated with the theoretical relative frequencies which would be expected from the schedules. This is because the points plotted represent the relative reinforcement values that actually occurred. Slight vagaries in the subjects' pattern of response and the variability inherent in the variable interval schedules account for this discrepancy. Had the theoretical values of the relative frequency of reinforcement been used in plotting these points, they would have been brought closer to the line representing equality of the percent response and reinforcement. However, it is customary to consider the schedules as they actually contact the subject (see, for example, Herrnstein, 1961), and this seemed more appropriate.

The dashed line represents the straight line fit to the data by the method of least squares. The slope of this line is .76 and the intercept is + 12.4.

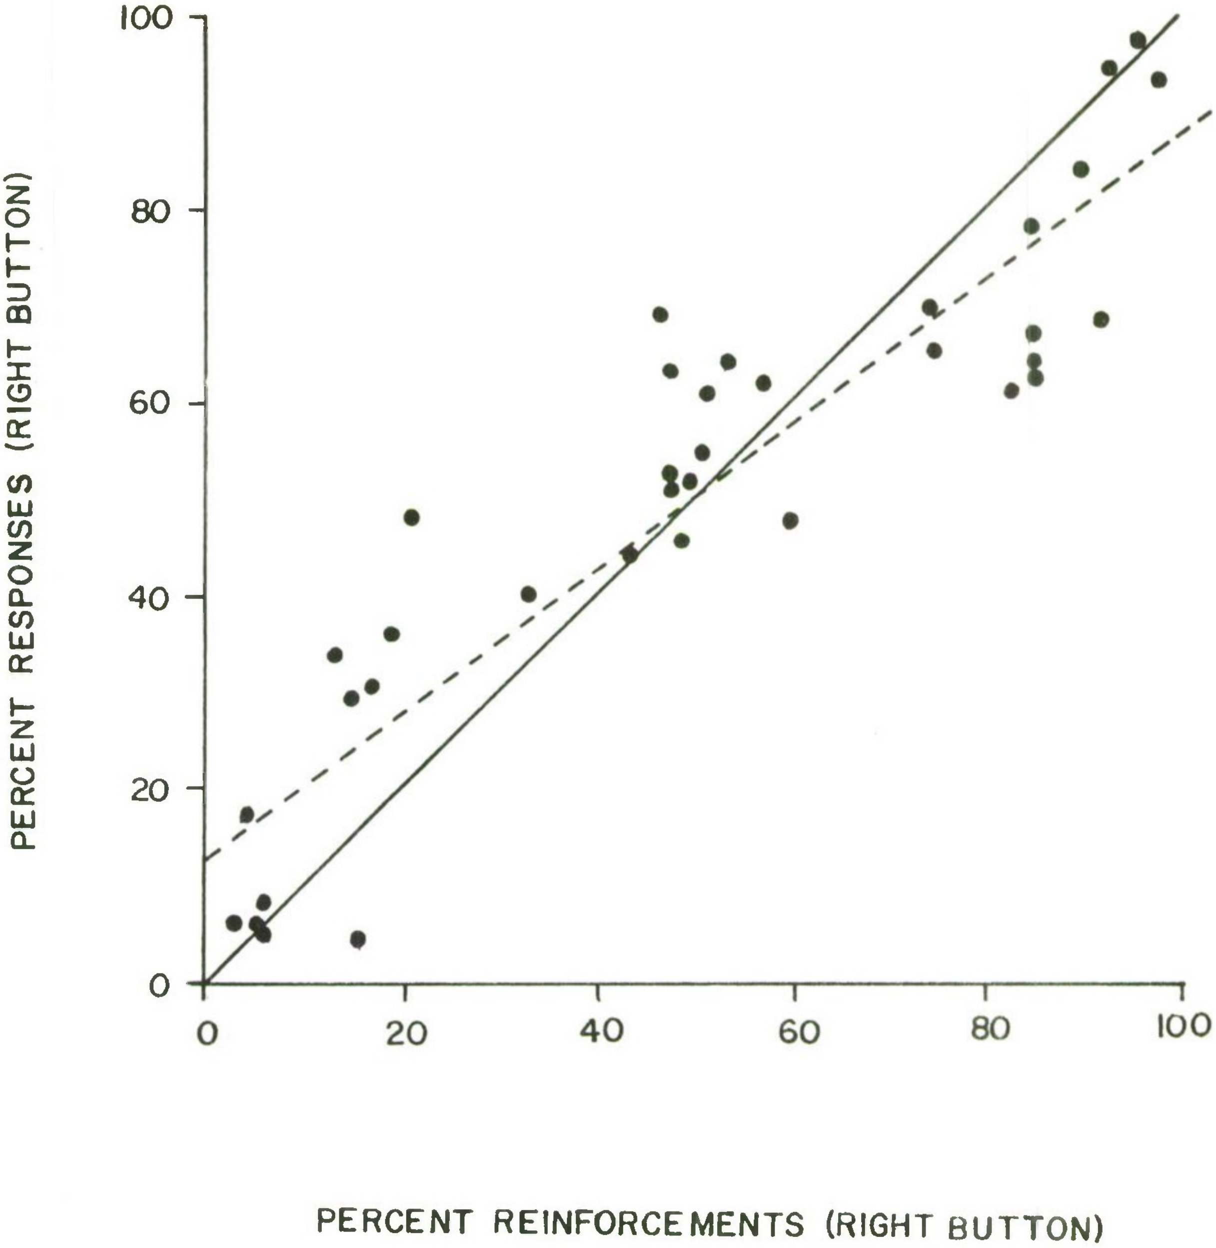

Fig. 4. Observed relationship between percent responses on the right button as a function of percent reinforcements on that button. Data for all fourteen subjects on concur. VI drl, VI drl schedules with several different mean interreinforcement intervals.

The discrepancy of these points from the expected values is again in the direction of overresponding on the button with the lower frequency of reinforcement.

Figure 5 shows the individual performances of four subjects. These data are plotted in the same manner as those in Fig. 4. Of all the subjects studied, subject CW (upper left) deviated least from the expected value. The largest discrepancy observed was for subject FB (lower right). In general, the data of the individual subjects closely parallels the group data.

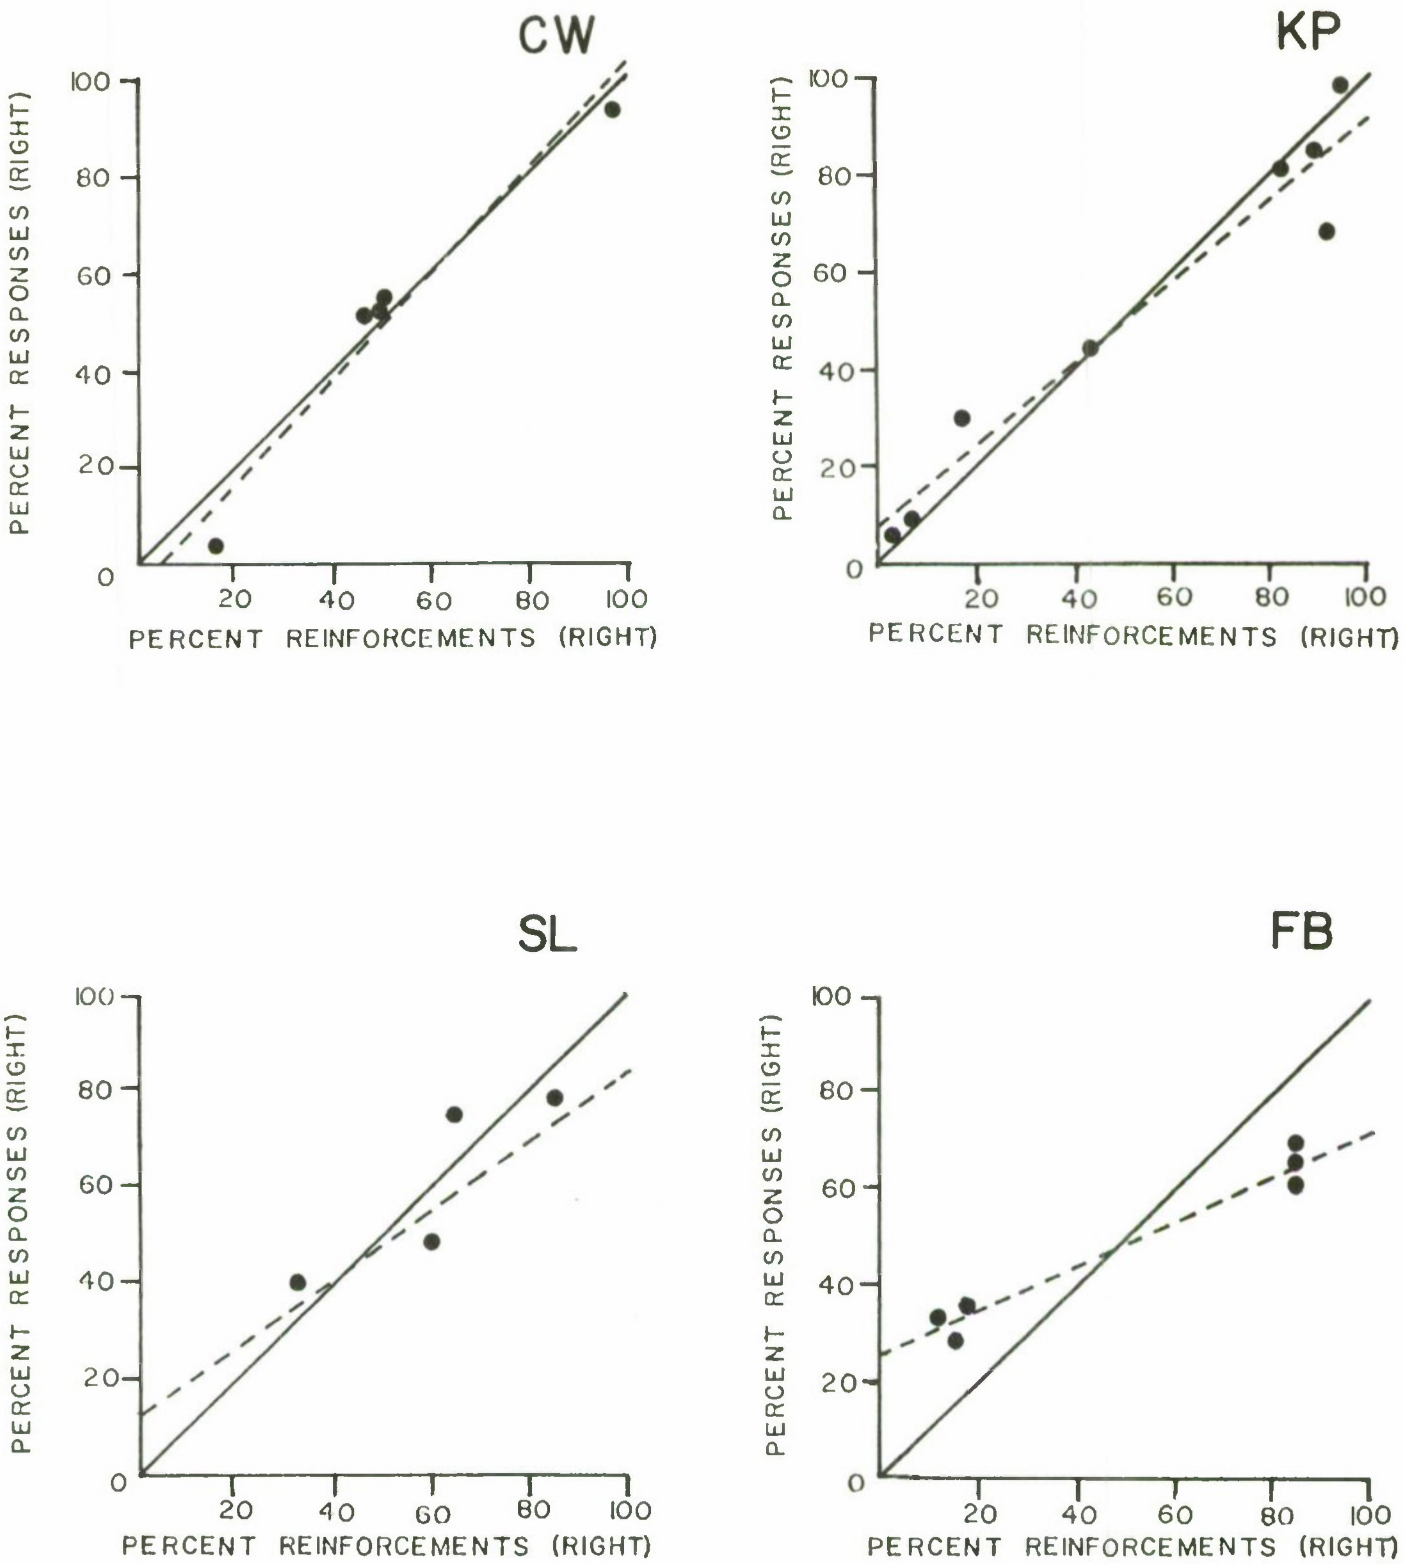

Fig. 5. Percent responses on the right button as a function of percent reinforcements on that button. Data for four individual subjects on concur. VI drl, VI drl schedules with several different mean interreinforcement intervals.

SECTION 4

DISCUSSION

The major outline of these results is clear: when two responses are under the control of concurrent reinforcement, these responses occur in proportion to the rates with which they are reinforced. Rate of response approximates a direct function of rate of reinforcement.

Other experiments with humans which studied only a single response (e.g., Hutchinson & Azrin, 1961; Holland, 1958) did not report such a direct relationship between rate of response and rate of reinforcement. Although these experiments were not basically concerned with this relationship, inspection of their data reveals that a linear relation was clearly not present. Only when two reinforced alternatives are available, as in the decision making paradigm, does this relationship emerge. Thus, these data complement the experiments with infrahumans. As Herrnstein pointed out in his experiments (1961), the linear relationship appears only when a second reinforced response is available.

The generality of these findings across species suggests that experiments in the animal laboratories will take on new importance for research in decision making. With the advent of studies of complex schedules of reinforcement they approach more closely the type of situations considered important for this research. The equipment and procedures for animal experiments have been developed to a high degree, and their suitability, particularly in terms of convenience, for extended periods of investigation make lower organisms valuable subjects for preliminary analysis. Because of such considerations, animal experiments may come to lead in this research.

The free operant response also has certain advantages. Absolute measures of rate can be studied as well as the relative measures employed in the present experiments. The study concerning the control exerted by rate of reinforcement per unit time mentioned in the introduction (Herrnstein, 1964b), shows how absolute measures of strength can be used not only to complement the findings based upon trial procedures, but also to extend these findings in new ways. Another advantage lies in the larger number of responses which can be observed using the free operant. Since large amounts of behavior can be observed in an individual organism, individual organism research becomes possible. The number of responses can replace the number of subjects in the statistical designs.

At first analysis, the linear relationship between rate of reinforcement and rate of response may appear to be simply a different way of wording the probability matching theorem common in decision theory. This theorem essentially states that the "probability of choosing a given alternative tends to match its probability of reinforcement" (Estes, 1962, p.428). The typical mathematical definition, however, bases the estimates of probability on number rather than on time. It is on this point that the two statements are discrepant, since time is the essential variable in "rate of reinforcement." For example, one alternative might be reinforced 1 time in 10 and another 1 time in 100, according to ratio schedules. By the probability matching theorem, we would expect that the responses would be distributed in proportion 10 to 1. But by such a distribution of responses (i.e., 10 to 1) the rate of reinforcement

would be less than if all responses were localized to the alternative reinforced 1 time in 10. Since it takes 100 responses on the other alternative to produce one reinforcement, 10 reinforcements could be obtained in that time if the subject remained on the key with the lower ratio. Thus, according to the finding reported here we would expect responses to occur exclusively on the one alternative.

Much of the experimentation in decision theory supports the probability matching theorem, but preliminary experiments with animals suggest that with extended exposure, responses tend to occur exclusively on the key with the lower ratio. The probability matching theorem may only characterize initial performance. When the difference in the rates of reinforcement for the two alternatives are slight, long periods of exposure may be necessary to see the effect. Additionally, superstitious chaining is likely to occur if a changeover delay is not provided. Thus, changes in the basic theorems of decision theory may result as these new findings are further explicated.

One of the problems raised by the present experiments was why the basic variable interval schedules, themselves, did not produce the expected relationship. A plausible explanation is simply that the behavior was not exposed to the schedule contingencies for a sufficient period of time. As has been pointed out, considerably more time was allocated for stabilization in the previous experiments upon which our expectations were based. And, in fact, there was an observed drift toward the expected values with longer periods of exposure. Practical considerations necessitated finding a procedure which produced stable performance more rapidly.

On the other hand, though, why was the drl contingency effective in speeding the acquisition of the expected relationship? At this point, no answers besides speculative ones can be given. It may be pointed out that a number of procedural factors, which superficially appear to be trivial, present similar problems of interpretation. The greater consistency found in human performance with fixed interval schedules when the force requirement for the response is increased, or when a penalty for responses is introduced, are examples. Even with experiments using infrahuman subjects we find similar problems. For example, simply adding a changeover delay requirement (Herrnstein, 1961) brings consistency where previously there was none. These apparently minor, nuisance considerations may in fact contain the answers to the important problems in predicting behavior. But only through further research can we expect that their importance will be drawn out and generalizations established.

REFERENCES

Azrin, N. H. Some effects of noise on human behavior. J. exp. Anal. Behav., 1958, 1, 183-200.

Catania, A.C., & Cutts, D. Experimental control of superstitious responding in humans. J. exp. Anal. Behav., 1963, 6, 203-208.

Estes, W.K. A descriptive approach to the dynamics of choice behavior. In E. Nagel, P. Suppes & A. Tarski (Eds.), Method and Philosophy of Science. Stanford: Stanford Univer. Press, 1962.

Ferster, C.B. & Skinner, B.F. Schedules of reinforcement. New York: Appleton-Century-Crofts, 1957.

Herrnstein, R.J. Relative and absolute strength of response as a function of frequency of reinforcement. J. exp. Anal. Behav., 1961, 4, 267-272.

Herrnstein, R.J. Secondary reinforcement and rate of primary reinforcement. J. exp. Anal. Behav., 1964, 7, 27-36 (a).

Herrnstein, R.J. Aperiodicity as a factor in choice. J. exp. Anal. Behav., 1964, 7, 179-182 (b).

Holland, J.G. Human vigilance. Science, 1958, 128, 61-63.

Hutchinson, R.R. & Azrin, N.H. Conditioning of mental hospital patients to fixed-ratio schedules of reinforcement. J. exp. Anal. Behav., 1961, 4, 87-95.

Skinner, B.F. Behavior of organisms. New York: Appleton-Century-Crofts, 1938.

Weiner, H. Some effects of response cost upon human operant behavior. J. exp. Anal. Behav., 1962, 5, 201-208.

Weiner, H. Conditioning history and human fixed-interval performance. J. exp. Anal. Behav., 1964, 7, 383-385.

Security Classification

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY *(Corporate author)*

Harvard University Committeeon Programmed Instruction
Cambridge, Mass.

2a. REPORT SECURITY CLASSIFICATION

UNCLASSIFIED

2b. GROUP

n/a

3. REPORT TITLE

Decision Making With Free Operant Responses

4. DESCRIPTIVE NOTES *(Type of report and inclusive dates)*

5. AUTHOR(S) *(Last name, first name, initial)*

Holz, William C.

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
June 64	22	11

8a. CONTRACT OR GRANT NO.

AF19(628)2404

b. PROJECT NO.

c.

d.

9a. ORIGINATOR'S REPORT NUMBER(S)

ESD-TDR-64-449

9b. OTHER REPORT NO(S) *(Any other numbers that may be assigned this report)*

10. AVAILABILITY/LIMITATION NOTICES

Qualified Requesters May Obtain from DDC.
Aval from OTS.

11. SUPPLEMENTARY NOTES

12. SPONSORING MILITARY ACTIVITY

Decision Sciences Lab,. ESD, L.G. Hanscom Field, Bedford, Mass.

13. ABSTRACT

These experiments explored the suitability of free operant techniques in the investigation of choice behavior and decision making. Young adults were the subjects. Two response manipulanda were available; and points were intermittently scheduled in different proportions for each. The number of points at the end of the session determined the subjects' payment. The schedule by which the points could result was the independent variable; and the relative frequency of the two responses, which represented the subject's choice, was the dependent variable. When the points were scheduled randomly in time, the anticipated result on the basis of previous findings was that the relative frequency of response would match the relative frequency of points. Theobserved result did not clearly follow this pattern. Over the period studied, the pattern was one of approximately equal responding to both choices regardless of the relative frequency of points obtained.

DD FORM 1 JAN 64 1473

Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Decision Making Behavior Reaction (Psychology) Human Engineering Learning Conditioned Reflex Experimental Data Programmed Instruction						

INSTRUCTIONS

1. ORIGINATING ACTIVITY: Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization *(corporate author)* issuing the report.

2*a*. REPORT SECURITY CLASSIFICATION: Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2*b*. GROUP: Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. DESCRIPTIVE NOTES: If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. AUTHOR(S): Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. REPORT DATE: Enter the date of the report as day, month, year; or month, year. If more than one date appears on the report, use date of publication.

7*a*. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7*b*. NUMBER OF REFERENCES: Enter the total number of references cited in the report.

8*a*. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report was written.

8*b*, 8*c*, & 8*d*. PROJECT NUMBER: Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9*a*. ORIGINATOR'S REPORT NUMBER(S): Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9*b*. OTHER REPORT NUMBER(S): If the report has been assigned any other report numbers *(either by the originator or by the sponsor)*, also enter this number(s).

10. AVAILABILITY/LIMITATION NOTICES: Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

(1) "Qualified requesters may obtain copies of this report from DDC."

(2) "Foreign announcement and dissemination of this report by DDC is not authorized."

(3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through ______________________."

(4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through ______________________."

(5) "All distribution of this report is controlled. Qualified DDC users shall request through ______________________."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. SUPPLEMENTARY NOTES: Use for additional explanatory notes.

12. SPONSORING MILITARY ACTIVITY: Enter the name of the departmental project office or laboratory sponsoring *(paying for)* the research and development. Include address.

13. ABSTRACT: Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as *(TS), (S), (C), or (U)*.

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. KEY WORDS: Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, rules, and weights is optional.